

Abstract

[0001] A technique is disclosed for queuing frames that are to be transmitted. The illustrative embodiment enables partial queuing of portions of frames on the part of a receiving entity, requiring only a minimum of real-time response on the part of the sending entity once the start of a frame is queued by the receiving entity. The illustrative embodiment relaxes queue memory requirements by recognizing that the interface between the sending entity and the receiving entity is able to respond to a request from the receiving entity to fetch additional frame data before the data presently in the partial queue is consumed. The amount of data that needs to be stored in each partial queue entry is causally related to a bounded latency time of the upper-to-lower layer interface multiplied by the physical layer rate of the transmitted frame.